

**ABSTRACT**

This disclosure describes leveraging workload information associated with executed database queries for estimating the result of a current database query. The workload information is analyzed to determine the usage of tuples in a database during query execution, such as how often a tuple is accessed and the number of different queries that accessed the tuple. A tuple is assigned a weight value that is based on the analyzed workload information. The particular tuples sampled for estimating a result for the current query is based on each tuple's weight value. The workload information may also be leveraged to generate an outlier index that identifies outlier tuples associated with the executed queries or that identifies outlier tuples associated with particular queries that are executed more frequently than other queries. The result for the current query can also be estimated using the sampled values along with the outlier tuples from the outlier index.